



US006041039A

United States Patent [19]
Kilkki et al.

[11] **Patent Number:** **6,041,039**
 [45] **Date of Patent:** **Mar. 21, 2000**

[54] **SYSTEM AND METHOD FOR DETERMINING NETWORK BANDWIDTH AVAILABILITY USING PRIORITY LEVEL FEEDBACK**

[75] Inventors: **Matti Kalevi Kilkki**, Espoo; **Jussi Pekka Olavi Ruutu**, Helsinki, both of Finland

[73] Assignee: **Nokia Telecommunications, OY**, Espoo, Finland

[21] Appl. No.: **08/821,222**

[22] Filed: **Mar. 20, 1997**

[51] Int. Cl.⁷ **G06F 11/00; H04L 12/56**

[52] U.S. Cl. **370/230; 370/235; 370/252; 370/395**

[58] Field of Search **370/230, 232, 370/233, 234, 235, 236, 252, 253, 395, 229, 468, 465, 231, 412**

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,942,569	7/1990	Maeno	370/230
4,953,157	8/1990	Franklin et al.	370/230
5,042,029	8/1991	Hayakawa	370/231
5,311,513	5/1994	Ahmadi et al.	370/468
5,539,747	7/1996	Ito et al.	370/235
5,546,389	8/1996	Wippenbeck et al.	370/412
5,724,451	3/1998	Chao et al.	370/395
5,828,653	10/1998	Goss	370/230

OTHER PUBLICATIONS

Technical Committee, Traffic Management Specification Version 4.0, af-tm-00561.000, Apr. 1996, pp. 1-59.
 Integrated Services Digital Network (ISDN) Overall Network Aspects and Functions, Traffic Control and Congestion Control in B-ISDN, ITU-T Recommendation, I.371, Mar. 1993, pp. 1-27.

What are the meaning of CBR, VBR, ABR, UBR?, Maintained by Carl Symborski, Last Changed Aug. 20, 1996, pp. 1-2 and 1-5.

ATM Service Categories: The Benefits to the User, Editor: Libio Lambaarelli, CSELT' Torino, Italy, pp. 1-10.

An Overview of ATM Technology, Gary Kessler, Jan. 1995, pp. 1-10.

ATM—The New Technology for Tomorrow's B-ISDN, Engui Yao, Dec. 1994, pp. 1-23.

Primary Examiner—Alpus H. Hsu

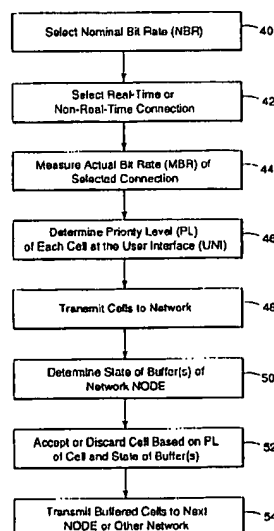
Assistant Examiner—Afsar M. Qureshi

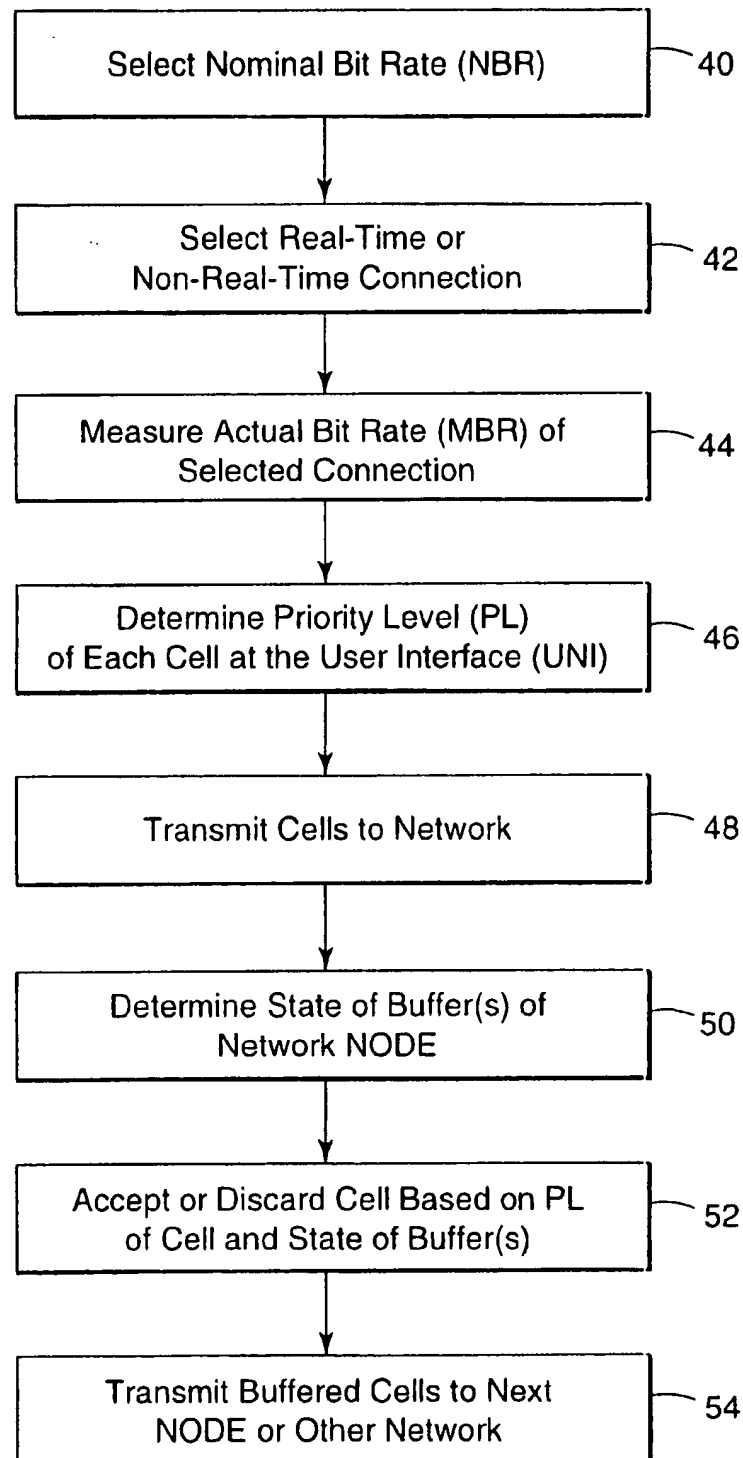
Attorney, Agent, or Firm—Altera Law Group, LLC

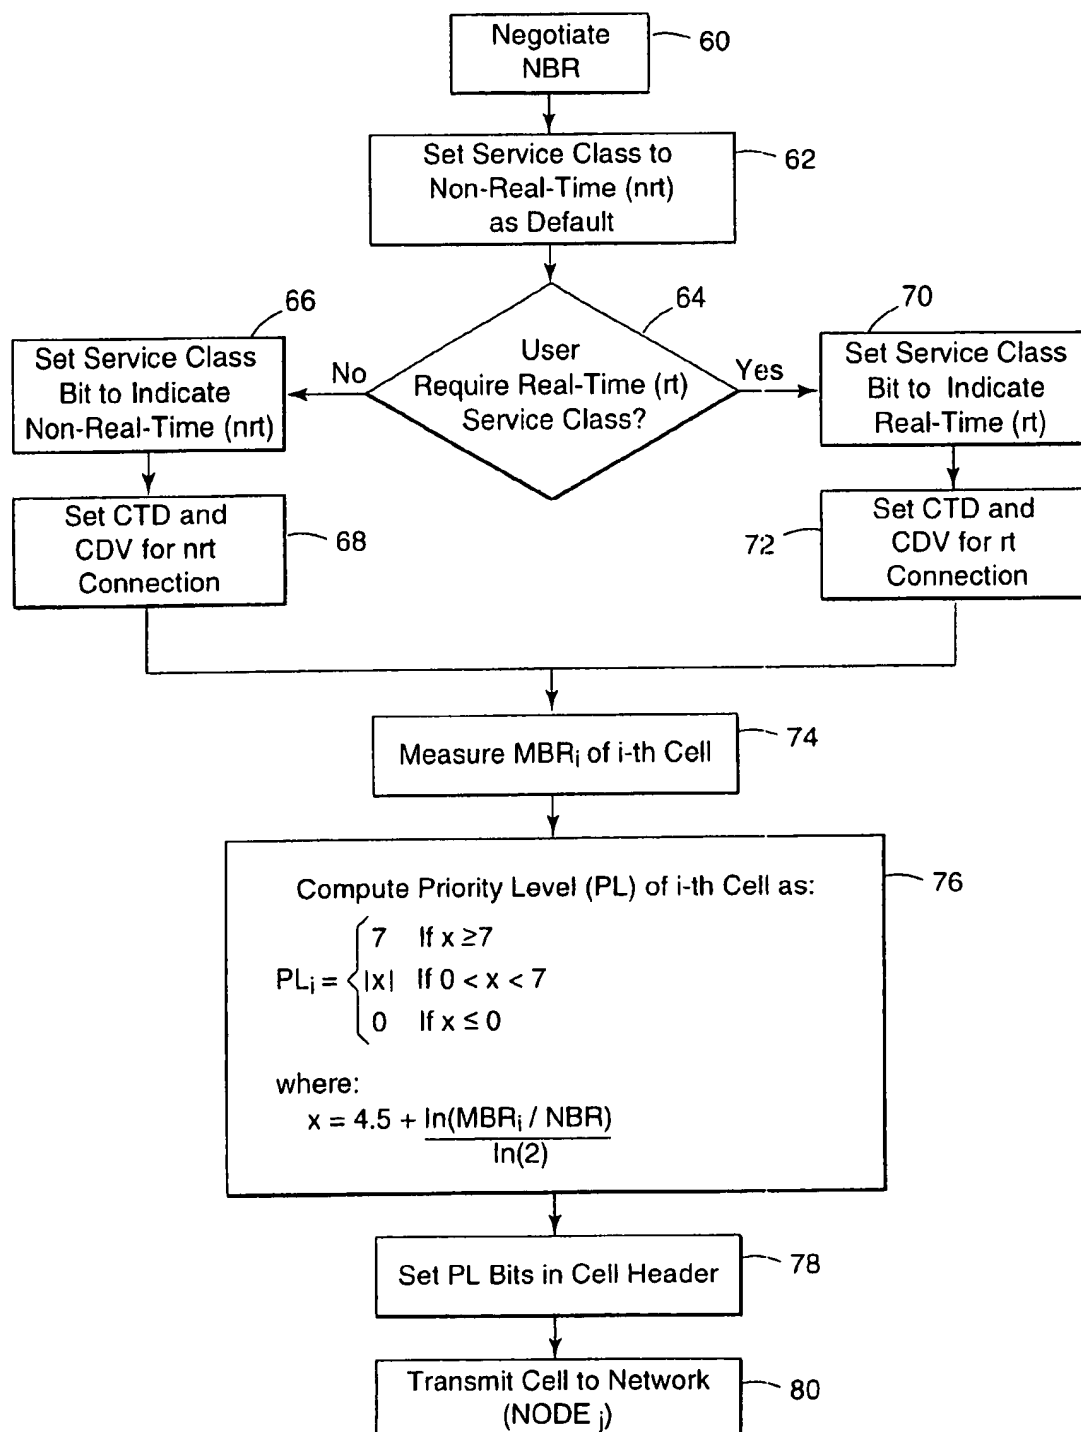
[57] **ABSTRACT**

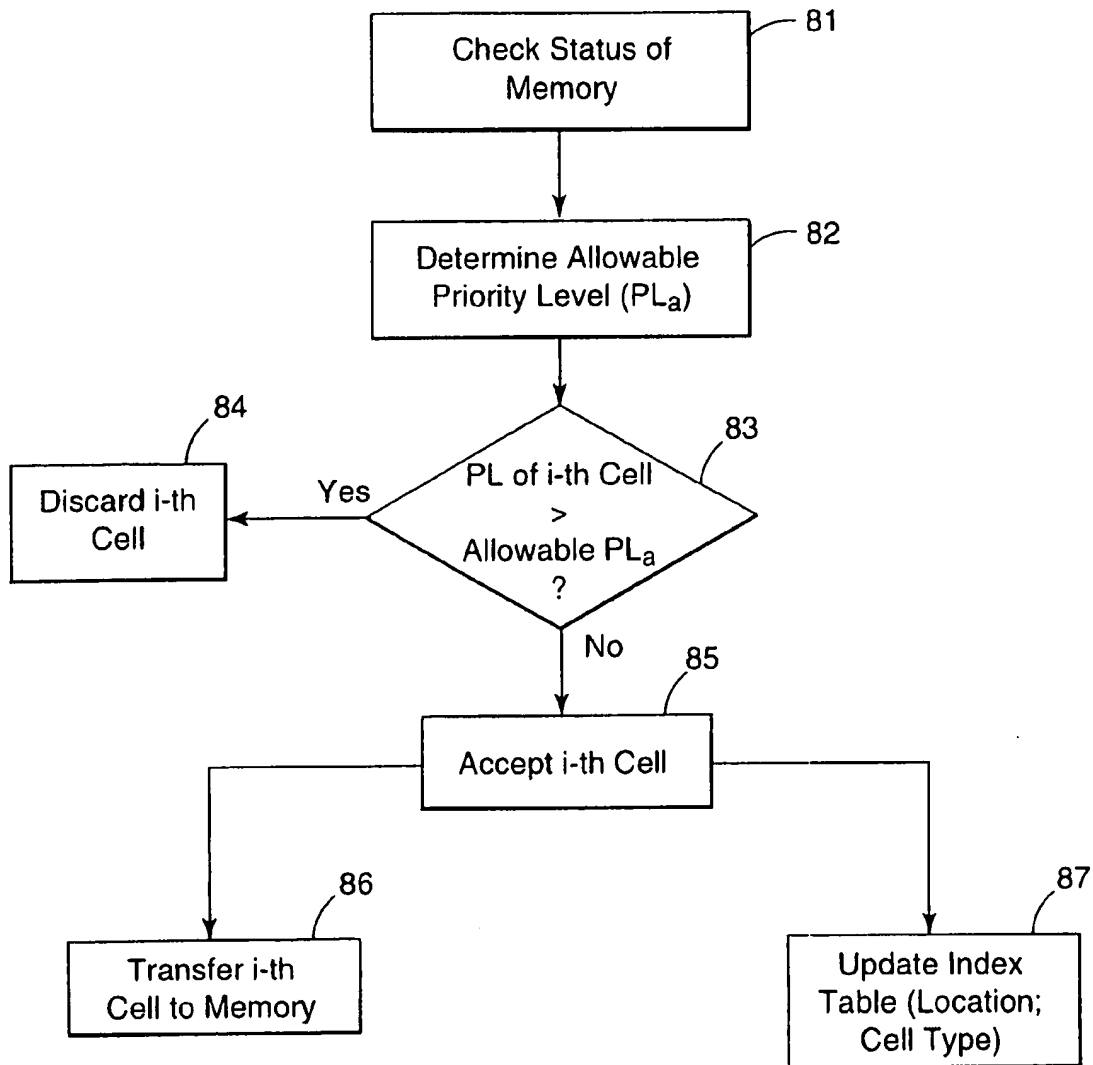
A system and method for managing information transfers over a network through priority level feedback is provided. The feedback feature provides priority level feedback which informs a cell source of a typical priority level which is being accepted by the nodes of the connection rather than being discarded due to connection congestion. The worst-case network load information for a connection is determined by calculating a requisite cell priority necessary to allow cell acceptance at each of the intermediary network nodes on the connection. A record is maintained of the minimum requisite cell priority corresponding to the worst-case network load information for the connection. The worst-case network load information is provided to the cell source, where one of the priority levels is established for cells subsequently output from the cell source in response to the network load information.

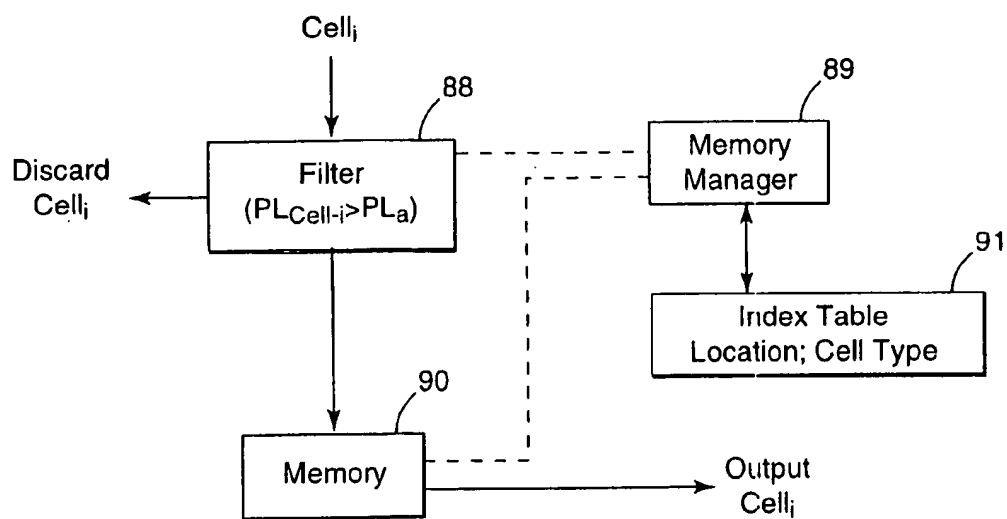
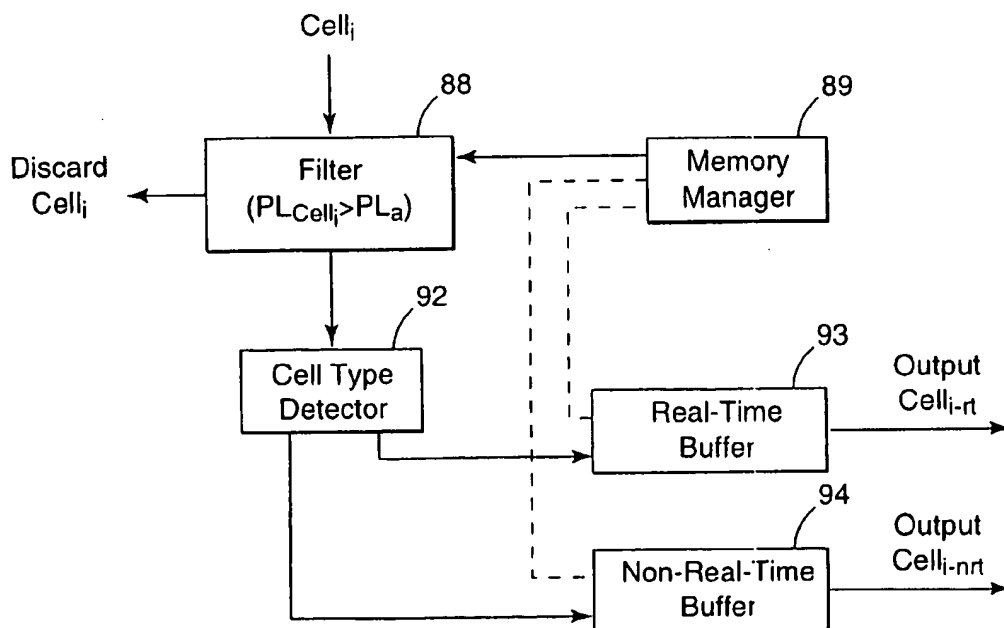
19 Claims, 10 Drawing Sheets



*Fig. 1*

*Fig. 2*

*Fig. 3*

*Fig. 4**Fig. 5*




Fig. 6




Fig. 7


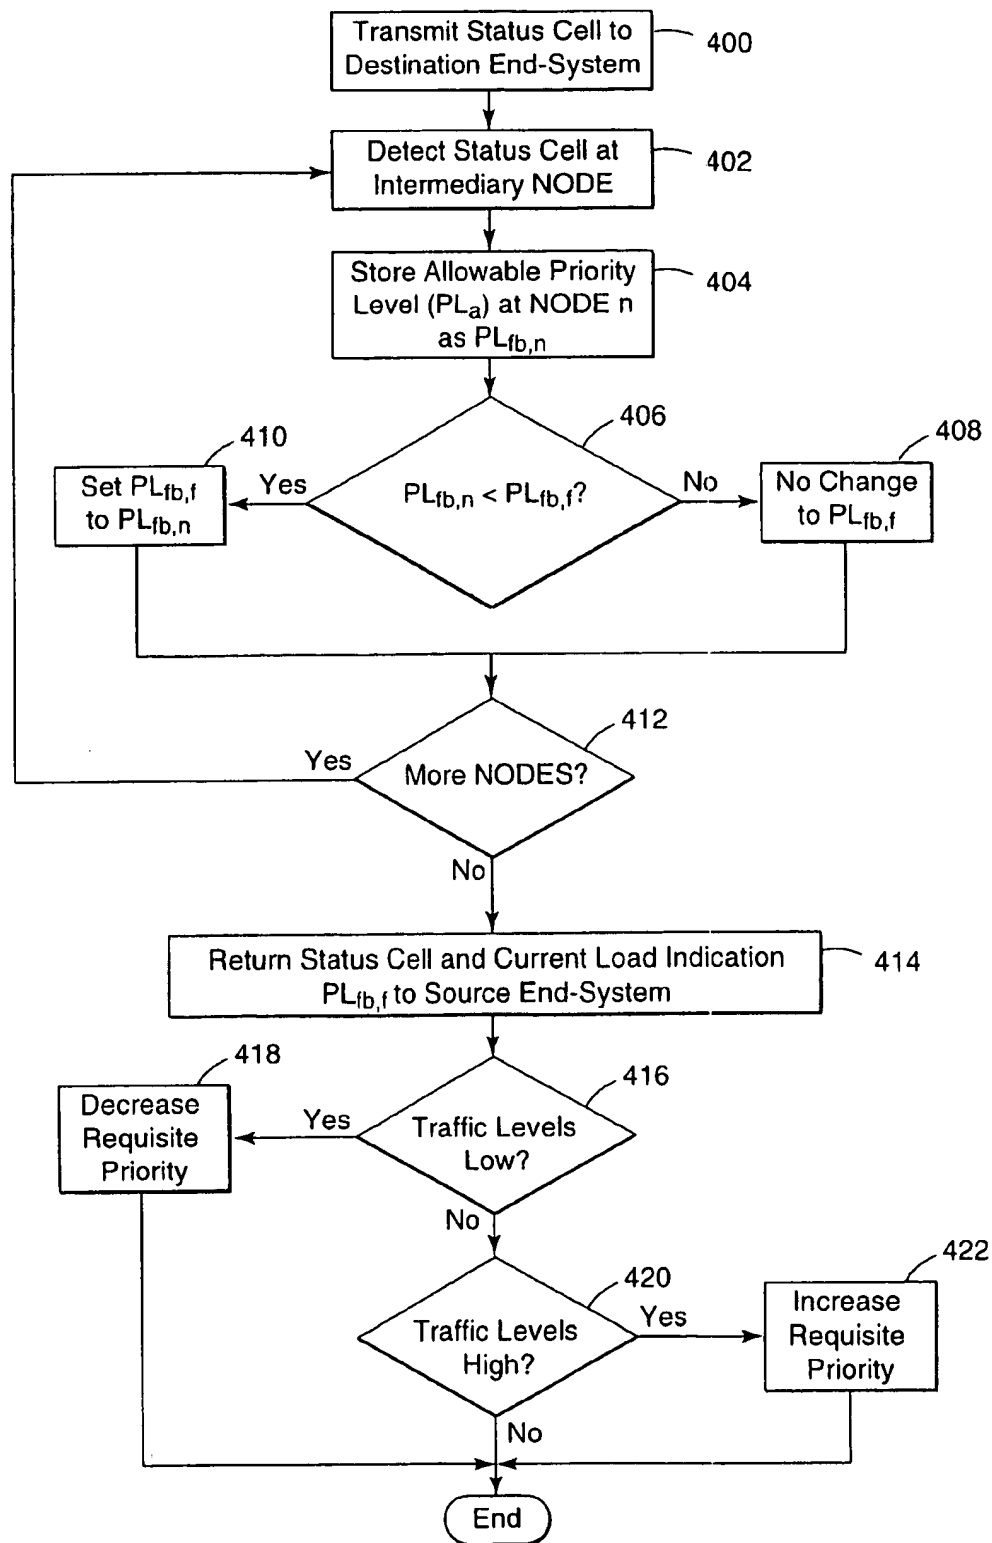










Fig. 8

*Fig. 9*



*Fig. 14**Fig. 15*

*Fig. 16**Fig. 17*

SYSTEM AND METHOD FOR DETERMINING NETWORK BANDWIDTH AVAILABILITY USING PRIORITY LEVEL FEEDBACK

FIELD OF THE INVENTION

The present invention relates generally to communications systems, and more particularly, to a system and method for managing information transfers over a network through priority level feedback.

BACKGROUND OF THE INVENTION

The state of communications technology, particularly that which affects the Internet, is currently in flux and subject to rapid and often uncoordinated growth. The ubiquity and variety of personal computers and set-top boxes has placed significant pressure on the providers of communications system infrastructure to accommodate the alarming increase in the number of new users that demand immediate access to Internet and other network resources. The rapid development of new and sophisticated software made available to users of such services places additional demands on system infrastructure.

Conducting commerce over the Internet and other networks is a practice that is gaining acceptance and popularity. By way of example, traditional on-line services, such as those offered by Internet providers, typical charge customers a monthly fee for access to basic services and resources, such as proprietary and public databases of information. Such traditional service providers also advertise any number of products or services which are purchasable on-line by the user.

Other forms of Internet commercialization currently being considered or implemented include offering of video and audio conferencing services, and a variety of other real-time and non-real-time services. The providers of these services, as well as the providers of communications system infrastructure, are currently facing a number of complex issues, including management of network capacity, load, and traffic to support real-time, non-real-time, and high-bandwidth services, and implementing a viable billing scheme that accounts for the use of such services.

The communications industry is expending considerable attention and investment on one particular technology, referred to as asynchronous transfer mode (ATM), as a possible solution to current and anticipated infrastructure limitations. Those skilled in the art understand ATM to constitute a communications networking concept that, in theory, addresses many of the aforementioned concerns, such as by providing a capability to manage increases in network load, supporting both real-time and non-real-time applications, and offering, in certain circumstances, a guaranteed level of service quality.

A conventional ATM service architecture typically provides a number of predefined quality of service classes, often referred to as service categories. Each of the service categories includes a number of quality of service (QoS) parameters which define the nature of the respective service category. In other words, a specified service category provides performance to an ATM virtual connection (VCC or VPC) in a manner specified by a subset of the ATM performance parameters. The service categories defined in the ATM Forum specification referenced hereinbelow include, for example, a constant bit rate (CBR) category, a real-time variable bit rate (rt-VBR) category, a non-real-time variable bit rate (nrt-VBR) category, an unspecified bit rate (UBR) category, and an available bit rate (ABR) category.

The constant bit rate service class is intended to support real-time applications that require a fixed quantity of bandwidth during the existence of the connection. A particular quality of service is negotiated to provide the CBR service, where the QoS parameters include characterization of the peak cell rate (PCR), the cell loss rate (CLR), the cell transfer delay (CTD), and the cell delay variation (CDV). Conventional ATM traffic management schemes guarantee that the user-contracted QoS is maintained in order to support, for example, real-time applications, such as circuit emulation and voice/video applications, which require tightly constrained delay variations.

The non-real-time VBR service class is intended to support non-real-time applications, where the resulting network traffic can be characterized as having frequent data bursts. Similarly, the real-time variable bit rate service category may be used to support "bursty" network traffic conditions. The rt-VBR service category differs from the nrt-VBR service category in that the former is intended to support real-time applications, such as voice and video applications. Both the real-time and non-real-time VBR service categories are characterized in terms of a peak cell rate (PCR), a sustainable cell rate (SCR), and a maximum burst size (MBS).

The unspecified bit rate (UBR) service category is often regarded as a "best effort service," in that it does not specify traffic-related service guarantees. As such, the UBR service category is intended to support non-real-time applications, including traditional computer communications applications such as file transfers and e-mail.

The available bit rate (ABR) service category provides for the allocation of available bandwidth to users by controlling the rate of traffic through use of a feedback mechanism. The feedback mechanism permits cell transmission rates to be varied in an effort to control or avoid traffic congestion, and to more effectively utilize available bandwidth. A resource management (RM) cell precedes the transmission of data cells, which is transmitted from source to destination and back to the source, in order to provide traffic information to the source.

Although the current ATM service architecture described above would appear to provide, at least at a conceptual level, viable solutions to the many problems facing the communications industry, ATM, as currently defined, requires implementation of a complex traffic management scheme in order to meet the objectives articulated in the various ATM specifications and recommendations currently being considered. In order to effectively manage traffic flow in a network, conventional ATM traffic management schemes must assess a prodigious number of traffic condition indicators, including service class parameters, traffic parameters, quality of service parameters and the like. A non-exhaustive listing of such parameters and other ATM traffic management considerations is provided in ITU-T Recommendation I.371, entitled Traffic Control and Congestion Control in B-ISDN, and in Traffic Management Specification, version 4.0 (atm-0056.000, April 1996), published by the Technical Committee of the ATM Forum.

One significant network traffic consideration is the amount of currently available bandwidth on a particular virtual connection. Except for the ABR service category, the existing ATM service classes do not suggest utilizing this network load information. The ABR service category provides for dynamic allocation of available bandwidth in response to network load information returned to the user.

However, the ABR service class provides a complex arrangement of feedback information, including a current

cell rate, explicit rate, minimum cell rate, direction indication, congestion indication, as well as others. This complex arrangement increases the complexity of the service class architecture. Also, the feedback mechanism defined for the ABR service class provides for allocation of bandwidth bounded between a defined minimum cell rate and a peak cell rate. Therefore, cell rate guarantees continue to exist which increases the complexity of the traffic management scheme. Furthermore, the conventional ATM service classes, including the ABR service class, do not provide solutions for determining network load conditions, and managing cell transfer rates in accordance with those network load conditions.

Accordingly, there is a need in the communications industry for a network arrangement and method that is not conceptually complicated to implement. A further need exists for a system and methodology that provides network capacity division during bandwidth overload situations, while providing network load information to the user with minimal complexity and negligible delays. The present invention fulfills these and other needs, and offers other advantages over the prior art traffic management approaches.

SUMMARY OF THE INVENTION

The present invention is directed to a system and method for managing information transfers over a network through priority level feedback. Network load information is provided to a network information source having a quality of service based on a nominal bit rate (NBR) service category. A nominal bit rate represents an expected, but not guaranteed, bit rate associated with a particular user or connection. Information elements transmitted from the sourcing unit to a network node are each assigned one of several priority levels. The present invention provides priority feedback information in an NBR system to allow regulation of the priority levels, which in turn provides for control of information transfer rates and realization of acceptable data loss probabilities.

In accordance with one embodiment of the invention, the feedback feature provides priority level feedback which informs a cell source of a typical priority level which is being accepted by the nodes of the connection, rather than being discarded due to connection congestion. The worst-case network load information for a connection is determined by calculating a requisite cell priority necessary to allow cell acceptance at each of the intermediary network nodes on the connection. A record is maintained of the minimum requisite cell priority corresponding to the worst-case network load information for the connection. The worst-case network load information is provided to the cell source, where one of the priority levels is established for cells subsequently output from the cell source in response to the network load information.

In accordance with another embodiment of the invention, a status cell is sent through the network connection to the destination, and returned to the cell source with the network load information. An acceptable priority field in the status cell is updated to reflect the minimum of the requisite cell priorities of the nodes as the nodes are reached by the status cell. When the status cell reaches the destination, the value of the acceptable priority field is stored in the status cell to be returned to the cell source. The cell source can adjust the priority level of subsequently output cells in accordance with the network load information.

The above summary of the present invention is not intended to describe each illustrated embodiment or imple-

mentation of the present invention. This is the purpose of the figures and the associated discussion which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

Other aspects and advantages of the present invention will become apparent upon reading the following detailed description and upon reference to the drawings in which:

FIG. 1 is a flow diagram illustrating a general procedure for communicating cells of information between a user/network interface and a network using a nominal bit rate service in accordance with an embodiment of the present invention;

FIG. 2 illustrates in greater detail a procedure for transmitting cells of information between a user/network interface and a network using a nominal bit rate service in accordance with another embodiment of the present invention;

FIG. 3 illustrates in flow diagram form a general procedure for filtering cells at a network node in accordance with an embodiment of a nominal bit rate service;

FIG. 4 is a block diagram of an embodiment of a system for filtering cells at a network node in accordance a nominal bit rate service;

FIG. 5 is a block diagram of an alternative embodiment of a system for filtering cells at a network node in accordance with a nominal bit rate service;

FIG. 6 is a block diagram illustrating one embodiment of an ATM network implementing the NBR methodology;

FIG. 7 illustrates one embodiment of a network load status cell in accordance with the NBR service of the present invention;

FIG. 8 is a block diagram of a representative ATM node within an ATM network configured to implement the NBR methodology;

FIG. 9 illustrates in flow diagram form a general methodology by which NBR status cells can be used to provide feedback to a source end-station to optimize cell transfer rates in accordance with the present invention;

FIG. 10 illustrates one embodiment for determining the priority level feedback PL_b wherein the minimum PL_a value is sought over a predetermined observation time period t_o ;

FIG. 11 illustrates another embodiment for determining the priority level feedback PL_b wherein a typical minimum PL_a value is determined by determining a minimum PL_a value occurring a predetermined number of times during a predetermined time period t_o ;

FIG. 12 illustrates an embodiment for determining the priority level feedback PL_b wherein the minimum PL_a value is determined by averaging the PL_a values over a time period t_o ;

FIG. 13 illustrates yet another embodiment for determining the priority level feedback PL_b wherein the minimum PL_a value is determined by calculating a weighted average of PL_a over a time period t_o ;

FIGS. 14-15 are graphical depictions of the relationship of average cell loss ratio, P_{loss} , as a function of priority level for four specific load levels;

FIG. 16 is a graphical illustration of the results of a simulation which compares the constant bit rate connections to those using feedback sources; and

FIG. 17 is a graphical illustration of the throughput/capacity from each of the feedback sources as a function of time.

DETAILED DESCRIPTION OF THE ILLUSTRATED EMBODIMENTS

In the following description of the various embodiments, reference is made to the accompanying drawings which

form a part hereof, and in which is shown by way of illustration various embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized, and structural and functional modifications may be made without departing from the scope of the present invention.

The service concept described herein may be viewed as a Simple Integrated Media Access (SIMA) service model. The SIMA service model incorporates the basic properties of ATM with the addition of eight priority levels as defined within the context of a new service concept termed a nominal bit rate (NBR) service. In general, the NBR service provides for the simple and efficient division of network capacity amongst different connections and the charging of user fees for the use of such connections. A network that embraces a basic version of the SIMA service model does not need to perform many of the traditional and burdensome traffic management functions involving traffic descriptors, quality of service parameters, service classes, connection admission control (CAC), or usage parameter control (UPC).

All of these functions are effectively replaced by functions performed by two autonomous units: a measuring unit, provided at a user/network interface, and a cell scheduling and buffering unit, provided at a network node. The SIMA service concept, from a user's perspective, is simple and understandable, because there are no pre-defined traffic or quality parameters associated with each connection, and charging for connection usage is based solely on the value of NBR and the duration of the connection.

A typical implementation of a SIMA service utilizes two primary components: access nodes and core network nodes, which have fundamentally different functional responsibilities. For example, access nodes, which may be a user/network interface, perform the task of measuring traffic for every connection, whereas at the core network nodes, the traffic control functions do not need to know anything about the properties of individual connections.

The elegant simplicity of the SIMA service model offers obvious advantages to the manufacturers of infrastructure hardware and software. For example, ATM switches or crossconnects can be built using individual cell scheduling and buffering units, switching fabrics, and routing functions. By using ATM virtual paths or IP switching technology, the routing tasks may be reduced in complexity. In addition, packet discarding and priority feedback features may be included in the cell scheduling and buffering units without negatively impacting their automaticity. Also, simple implementation of network nodes may result in the availability of a relatively inexpensive, high capacity network infrastructure.

The more complex unit of the SIMA service infrastructure concerns the access nodes. Such access nodes will typically include a measuring unit to measure the traffic stream of every connection in real-time, and a computation unit for determining a priority to be assigned to every cell. These additional features should be implementable at a level of difficulty no greater than that for effecting UPC in conventional ATM networks.

The present invention provides priority feedback information in an NBR system as disclosed herein, so that cell-sourcing units are able to regulate their cell transfer rates (CTR) and realize acceptable cell loss probabilities. Some information regarding the connection threshold level is periodically updated and fed back to the cell-sending unit. Each ATM node calculates a typical allowable priority level

that can be inserted into special status cells used by the source to probe the connection.

In one embodiment of the invention, this feedback feature is used to optimize the CTR by maintaining the highest priority possible without exceeding the worst-case allowable priority accepted at the nodes of the virtual connection. The invention, therefore, provides priority level feedback (PL_{fb}) which informs the source end-system of a typical priority level that is still accepted by the nodes of a connection.

In order to provide the priority level feedback, the allowable priority level PL_a is determined at each node in the connection. A priority level feedback level PL_{fb} records the worst-case PL_a from the source to the destination, which is stored as a field of the status cell. The present invention includes various embodiments for determining the value of the lowest allowable priority (i.e., the highest allowable priority "level", PL_{fb}) currently being accepted on the connection.

Subsequent ATM nodes compare the PL_{fb} in the special status cell to the current priority level PL_a that is being accepted at the node (i.e., not discarded at the node). If the status cell includes a PL_{fb} having a larger value than the current allowable priority level at the node, then PL_{fb} in the status cell is replaced with a new, smaller value corresponding to the PL_a value at the node. The connection destination will thus receive the smallest PL_{fb} of the connection, indicating what is typically the highest priority level PL (corresponding to the worst priority) that will successfully traverse the connection without cell discarding. The destination unit then sends this network load information back to the cell source, which allows the user to adjust the CTR in order to reduce the likelihood that subsequently issued data cells will be lost through cell discarding at the nodes of the connection.

Referring now to FIG. 1, there is shown a general methodology for transmitting information between a user/network interface and a network over an NBR service connection. Initially, a user negotiates or selects a nominal bit rate with the network operator, which may be performed prior to, or at the time of, establishing the connection. In one embodiment, the user informs the network operator that a desired NBR is required, and the requested connection bandwidth is allocated to the user. The network operator, in accordance with this embodiment, need not perform the task of analyzing current network load conditions existing at the core network nodes prior to establishing or releasing an NBR connection. In an alternative embodiment, the network operator performs the task of determining network load status prior to establishing or releasing an NBR connection, although this task may not be necessary in a properly dimensioned network supporting an NBR service.

Depending on a particular application, the user selects a real-time or a non-real-time network connection. The process of determining the priority level (PL) of each cell, which indicates the importance or criticality of the cell relative to other cells, involves measuring the actual or measured bit rate (MBR) of the selected real-time or non-real-time connection at the UNI. The priority level of each cell is determined at the UNI using a ratio of MBR and NBR.

After computing the priority level of each cell at the UNI, the cells are transmitted to the network, such as to a node of the network. A network node, upon arrival of a cell transmitted from the UNI, performs a cell filtering process by which the node determines whether to accept or discard

a particular cell. The cell filtering process involves determining 50 the state of one or more buffers or memories of the network node to determine a buffer or memory occupancy level. The node accepts or discards 52 a cell based on the priority level of the cell and the state of the node buffer. Cells that meet the filtering criteria determined at the node are accepted, buffered, and eventually transmitted 54 to another node in the network or another network in a manner consistent with the expected quality of service for the connection.

FIGS. 2-4 illustrate a procedure for scheduling and buffering cells in accordance with one embodiment of an NBR service methodology. Referring now to FIG. 2, a user establishes 60 an NBR with a network operator. It may be desirable, although not required, to initially set the service class 62 to a non-real-time (nrt) service class as a default setting. Depending on a particular application, the user may require a real-time (rt) service class 64, which may be set by the user directly or, typically, by the user's application or communications software. If the user requires a real-time connection, each cell transmitted from the user's UNI will have the service class bit in the cell header set to indicate that the payload of the cell contains real-time information 70. It is noted that within the context of a network implemented in accordance with the NBR concept of the present invention, real-time service class connections are expected to support virtually any real-time application without the need to specify particular cell transfer delay (CTD) and cell delay variation (CDV) parameters. As such, the conventional procedure of setting CTD and CDV bits of the cell header to appropriate values to accommodate the real-time service requirements of the connection is altogether obviated.

If the user does not require a real-time service connection, the default non-real-time service class condition remains operative. As such, the rt/nrt service class bit of each cell header is set to indicate that the payload of the cell includes non-real-time information 66. It is noted that the NBR service disclosed herein does not utilize the cell loss priority (CLP) scheme used by conventional ATM traffic management approaches. As such, the CLP bit in the cell header may instead be used to discern between real-time and non-real-time payloads.

In the above described embodiment, each cell transmitted over a connection is designated as either a real-time cell or a non-real-time cell, such as by appropriately setting the rt/nrt service class bit of the cell header. In an alternative embodiment, depending on a user's requirements, a connection may be designated as being either a real-time or non-real-time connection, and the cells communicated over such a connection need not be individually assigned a real-time or non-real-time status. Each node for a given connection, for example, may perform a table look up procedure upon arrival of a cell at the node to determine whether the cell is associated with a real-time or a non-real-time connection. Thus, in accordance with this embodiment, a cell header bit need not be reserved for distinguishing between real-time and non-real-time cells.

After the rt/nrt service class header bit has been set in the above-described manner, the actual bit rate of a particular cell to be transmitted between the UNI and the network is measured 74. Since, in practice, the actual bit rate may be subject to significant variability over time, a measuring unit of the UNI employs an averaging measuring principle to determine the actual or instantaneous bit rate, MBR_i .

In general, the UNI measures 74 the actual bit rate of a cell, such as cell_i, by approximating the actual or instantane-

ous bit rate of the connection within a measuring period having a duration that is appropriate for the particular connection (e.g., a real-time or non-real-time connection). The instantaneous bit rate, MBR_i , may be determined using a known technique.

Having determined 74 the measured bit rate, MBR_i , of the i :th cell, the priority level of the i :th cell is computed using the measured bit rate, MBR_i , and the nominal bit rate, NBR. In accordance with one embodiment, it is assumed that a cell may be distinguished from other cells using a cell prioritization scheme that employs eight priority levels. In order to indicate which of the eight priority levels is attributed to a particular cell, each cell allocates three bits for this purpose.

In accordance with current ATM specifications, an ATM cell is specified as a unit of transmission having a fixed-size frame consisting of a 5-octet header and a 48-octet payload. It is appreciated that the necessity to allocate three bits in the cell header for the purpose of designating cell priority level may require utilization of currently defined ATM header bits. By way of example, it may be possible to utilize the current Generic Flow Control (GFC) field which constitutes a total of four bits. In this case, three bits may be allocated for designating cell priority level and one bit may be designated as the rt/nrt service class bit. It may be possible, in accordance with another embodiment, to allocate other header bits for the purpose of indicating one of eight priority levels and rt/nrt service class by deviating from the five-octet header ATM specification.

As such, other header bits may be redefined to represent cell priority level and service class designations. Alternatively, one or more bits required to specify cell priority level and/or service class may be situated outside of the currently defined ATM cell header. The need to make a minor modification to the existing ATM cell header definition is significantly offset by the substantial advantages offered by employing the NBR service scheme of the present invention, such as a significant reduction in network and traffic management overhead and complexity.

It is understood that the number of priority levels may be less than eight or greater than eight. By way of example, if it is assumed that four cell header bits are allocated for purposes of indicating a cell's priority level, as many as 2^4 (i.e., 2^{n-bits}) or 16 priority levels may be defined. Increasing the number of priority levels within the context of an NBR service permits the network operator to make finer adjustments to the bandwidth of a particular connection when managing network traffic. The price for this finer level of traffic control is the additional cell header bit or bits needed to resolve a greater number of priority levels.

A priority level computing unit determines 76 the priority level of each cell, such as cell_i, using the computed value of MBR_i and the value of NBR. In accordance with one embodiment of the present invention, and assuming that the measured bit rate is MBR_i when the i :th cell is transmitted to the network, the priority level (PL_i) of cell_i may be calculated using the following equation:

$$x = 4.5 + \frac{\ln(MBR_i / NBR)}{\ln(2)} \quad (1)$$

$$PL_i = \begin{cases} 7 & \text{if } x \geq 7 \\ \lfloor x \rfloor & \text{if } 0 < x < 7 \\ 0 & \text{if } x \leq 0 \end{cases}$$

where, $\lfloor x \rfloor$ represents the integer part of x . As will be discussed hereinbelow in accordance with an embodiment in

which both NBR and traditional ATM service connections are accommodated, the zero priority level, $PL=0$, is reserved for those connections that use an ordinary ATM service with guaranteed bandwidth and quality of service. Accordingly, Equation [1] above may be modified so as to produce cell priority levels ranging between $PL=1$ and $PL=7$, such that $[x]$ represents the integer part of x , if $1 < x < 7$.

It can be seen by application of Equation [1] above that if a connection is exploiting network capacity in excess to the connection's negotiated NBR value, the priority level of cell, is at least 4. It can further be seen that if the momentary bit rate at the UNI is less than the negotiated value of NBR, PL is at most 4. The priority level scheme in accordance with this embodiment of the present invention thus permits adjustment of the relative capacity used by a connection in steps of 2. From Equation [1] above, it can be seen that for an NBR of 100 kbit/s, an MBR higher than 566 kbit/sec results in a PL of 7, and an MBR lower than 8.8 kbit/s results in a PL of 0.

The three priority level bits allocated in the cell header are set 78 for each ATM cell transferred from the UNI. The ATM cells are then transmitted 80 to targeted network nodes, identified by node addressing information provided in the cell header.

It is noted that if a user is not satisfied with the quality of service of the connection, the user may elect between at least three alternatives. First, the user may elect to keep the average bit rate unchanging, but reduce the variation of traffic process. Second, the user may elect to decrease the average bit rate, or to increase the nominal bit rate. Increasing the NBR will, however, generally result in a concomitant increase in cost for a higher speed connection. Finally, the user may change the network operator.

In FIG. 3, there is illustrated in flow diagram form a general methodology by which a network node processes cells containing priority level information received from a UNI in accordance with one embodiment of the present invention. FIG. 4 illustrates an embodiment of various components of a network node employed to effectuate the methodology illustrated in FIG. 3. It is assumed that a cell, such as cell_i, has been processed at a UNI and includes priority level information derived in a manner described hereinabove.

Cell_i is transmitted from the UNI to a network node and is received at a filter 88 of the node. A memory manager 89 checks the status 81 of the memory 90 in order to determine the occupancy in the memory 90. The memory manager 89 determines 82 the allowable priority level (PL_a) based on the occupancy state of the memory 90. In general, the memory manager 89 establishes a high allowable priority which translates to a low allowable priority "level," for example $PL_a=0$ or 2, when the occupancy level of the memory 90 is high (i.e., few available memory locations). When the memory manager 89 determines that the memory 90 has ample capacity for receiving new cells, the memory manager 89 establishes a low allowable priority which translates to a high allowable priority "level," for example $PL_a=6$ or 7. As will be appreciated by those skilled in the art, the calculation of PL_a could alternatively be based on unoccupied buffer capacity rather than on buffer occupancy without departing from the spirit of the invention.

If the priority level of cell_i is greater than the allowable priority level, PL_a , as determined 83 by the memory manager 89, the filter 88 discards 84 cell_i. If, on the other hand, the priority level of cell_i is equal to or less than the allowable priority level PL_a , the filter 88 accepts 85 cell_i. The memory manager 89 coordinates the transfer 86 of cell_i to the

memory 90 and updates an index table 91 coupled to the memory manager 89 to include a new index table entry for newly accepted cell_i. In one embodiment, the index table 91 stores the location of the accepted cell, in the memory 90, and also stores a cell-type indicator which specifies whether cell_i is a real-time cell or a non-real-time cell. As such, the memory 90 may store both real-time and non-real-time cells.

The memory manager 89, in cooperation with the index table 91, manages cell transfer operations from the memory 90 to the output of the memory 90 by giving preference to the real-time cells over the non-real-time cells. By way of example, the memory manager 89, upon determining the presence of both rt-cells and nrt-cells stored in the memory 90, transfers all of the rt-cells to the output of the memory 90 prior to transferring out any of the nrt-cells.

In accordance with another embodiment, as is illustrated in FIG. 5, a memory manager 89 determines the status of a real-time buffer (rt-buffer) 93 and a non-real-time buffer (nrt-buffer) 94. The memory manager 89 determines, in a manner similar to that previously described, the allowable priority level, PL_a , for the filter 88 based on the status of the rt-buffer 93 and the nrt-buffer 94. If the priority level of cell_i is greater than the allowable priority level, PL_a , the filter 88 discards cell_i. If, on the other hand, the priority level of cell_i is equal to or less than the allowable priority level, PL_a , cell_i is accepted.

In accordance with another embodiment, the network node may apply a buffer filtering scheme which performs the filtering function based on packets of cells, rather than on individual cells. By way of example, the filtering procedure described hereinabove may be applied to the first cell of each packet. If the first cell is discarded by the node, then all of the cells of the packet following the first cell are discarded as well. If, however, the first cell of a packet is accepted, then the priority of all other cells belonging to that packet may be increased, for example by changing the priority level from $PL=5$ to $PL=3$. A gain of even one priority level, such as from $PL=4$ to $PL=3$, is believed to be sufficient to ensure that there will only be very few partially transmitted packets.

A cell-type detector 92 receives the accepted cell, cell_i, from the filter 88 and determines whether cell_i is an rt-cell or an nrt-cell. As discussed previously, the header of cell_i includes a header bit, such as the CLP bit, which indicates whether or not cell_i is an rt-cell or an nrt-cell. The cell-type detector 92, upon determining the service class type of the cell_i, transfers the cell_i to either the rt-buffer 93 or the nrt-buffer 94. In a manner similar to that described previously with respect to FIGS. 3 and 4, the memory manager 89 coordinates the output of rt-cells and nrt-cells respectively from the rt-buffer 93 and the nrt-buffer 94, giving preference to the rt-cells.

In accordance with another embodiment of the present invention, it may be desirable, for purposes of enhancing network expansion and traffic control, to request that each user of the network purchase a maximum NBR. The maximum NBR value is intended to remain substantially constant. In addition, it may be desirable to request that each user select an appropriate instantaneous NBR, which should be no greater than the selected maximum NBR. The selection of an appropriate instantaneous NBR generally involves a compromise between price and quality of service. The service quality detected by a user depends largely on three parameters, namely the NBR, the average bit rate, and the amount of traffic variations. Although a user may change any of these parameters, the only information that the network needs to know at the initiation of cell transmission is the NBR and the service class (real-time or non-real-time) of the connection.

In accordance with another embodiment of the present invention, a SIMA service model provides for the accommodation of both NBR and traditional ATM service connections. It is appreciated that traditional ATM services which offer guaranteed connections may be desirable for certain applications. It is anticipated, however, that the quality of service offered by the NBR service of the present invention will meet or exceed a user's expectations for virtually all real-time and non-real-time applications.

A SIMA service which provides for both NBR and traditional ATM services requires that the network operator dedicate a UPC device for each conventional ATM connection, or possibly for each virtual path. All of the cells transmitted using traditional ATM service connections are designated with the highest priority of $PL=0$ and with a real-time (rt) service class designation. In accordance with this approach, the zero priority level is reserved for those connections that use an ordinary ATM service with guaranteed bandwidth and quality of service. Accordingly, the priority determination Equation [1] above is modified so as to produce cell priority levels ranging between $PL=1$ and $PL=7$, such that $\lfloor x \rfloor$ represents the integer part of x , if $1 < x < 7$. It is noted that if the network operator wants to mark excessive cells as $CLP=1$ cells, those cells may be marked with a lower priority, such as $PL=6$, for example.

A point of possible incompatibility with traditional ATM technology involves the necessity of three bits for each ATM cell for the determination of cell priority, or two bits if the current cell loss priority, CLP, bit in the cell header is used. In addition, one bit is needed to distinguish between real-time and non-real-time connections. The rt/nrt service bit may be, but is not required to be, included in every cell. It may be possible to utilize the current Generic Flow Control (GFC) field which constitutes a total of four bits. In this case, three bits may be allocated for designating cell priority level and one bit may be designated as the rt/nrt service class bit.

FIG. 6 is a block diagram illustrating one embodiment of an ATM network 200 implementing the NBR methodology. The exemplary ATM network 200 used for purposes of this description is depicted as a network having two intermediary ATM nodes 202 and 204. However, as will be appreciated by those skilled in the art, the present invention may likewise be implemented in various network structures such as multipoint, star, ring, loop and mesh network topologies used in networks ranging from local area networks (LAN) to proliferative global area networks (GAN) such as the Internet.

The network 200 includes a source end-system 206 for transmitting digital information to the destination end-system 208. The information transmitted in such a network typically passes through various network nodes, such as nodes 202 and 204, as it advances towards its destination. These nodes represent network data communications elements such as routers, switches or multiplexers. Connecting the end-systems and nodes are circuit connections that provide the means for which the digital information may be passed. Connection links 210, 212 and 214 represent the connections for data sent from the source end-system 206 to the destination 208, while connection links 216, 218 and 220 represent those connections which provide information in return.

FIG. 6 further illustrates an ATM cell flow in the ATM network 200 implementing the NBR methodology. As data is sent to the destination end-system 208 by way of a flow of ATM cells along connections 210, 212 and 214, network load information may be returned to the source end-system 206 via connections 216, 218 and 220. The NBR system

operates on a priority basis, thereby resulting in the presentation of NBR network load information as priority level information. Information regarding current allowable priority levels at the nodes is provided to the source end-system 206 in order to provide status and allow optimization of the cell transfer rate (CTR).

In one embodiment of the invention, the network load information is provided to the source end-system 206 in the form of special ATM status cells periodically issued from the source end-system 206. The status cells are part of the normal connection cell flow in the sense that they are included in the normal MBR calculation of the connection, and the status cell priority level is calculated in the manner previously described in connection with FIG. 2. FIG. 6 illustrates the progress of a typical status cell in accordance with the present invention, seen at six different intervals as it advances from source end-system 206 to the destination end-system 208 and back again. The status cell 222a-f is depicted at times $t=1$ through $t=6$ respectively, which corresponds to the location/time relationship as the status cell travels along a connection.

Referring now to FIG. 7, one embodiment of a network load status cell 250 in accordance with the present invention is shown. ATM standards define an ATM cell as a fixed-size cell with a length of 53 octets comprised of a 5-octet header and a 48-octet payload. The status cell 250 is modeled after the ATM standard cell, and includes a 5-octet header 252 and a 48-octet payload 254. Within the payload section 254 of the status cell 250 are a pair of priority level feedback (PL_{fb}) cells, labeled $PL_{fb,f}$ 256 (forward priority level feedback) and $PL_{fb,b}$ 258 (backwards priority level feedback), that are used to store priority level information as the status cell 250 travels from source-to-destination and destination-to-source respectively. The $PL_{fb,f}$ field 256 gathers the worst-case PL_a on the connection identified by the value of the lowest allowable priority (i.e., the highest allowable priority "level", PL_{fb}) currently being accepted from the source to the destination. In one embodiment, the source end-system 206 initially sets the $PL_{fb,f}$ field 256 to the lowest priority, corresponding to a priority level value of "7".

As each node receives the status cell 250, it examines its current load level at the outgoing link of the connection. The load condition at a particular node is denoted $PL_{b,n}$ which identifies the highest priority, and therefore the lowest allowable priority level PL_a , for the ATM node. The current load condition $PL_{b,n}$ at the node is compared to the value available in the $PL_{fb,f}$ field 256, where the $PL_{fb,f}$ field 256 reflects the lowest allowable priority level PL_a , identified in a node on the connection. Where $PL_{b,n} < PL_{fb,f}$ the value in the $PL_{fb,f}$ field 256 is decreased to reflect the lowest allowable priority level identified up to this point in the connection, and is therefore decremented to equal the PL_a value of the node. Where $PL_{b,n} > PL_{fb,f}$ the node does not change the value in the $PL_{fb,f}$ field 256.

Each network node detects the status cell 250 based on identifying information in the cell header 252. Among the ATM header fields is a 3-bit payload type (PT) field 260 used to discriminate between a cell payload carrying user information from a cell payload carrying management information. The PT field 260 of the header 252 in FIG. 7 is used to distinguish a status cell 250 from a standard data cell. Any desired combination of bits in the PT field 260 could be used to identify a status cell 250. Alternatively, a separate bit in another location of the header 252 can be used to distinguish status cells 250 and standard data cells.

After the destination end-system 208 has received the status cell 250, it returns the status cell 250 to the source

end-system to allow it to examine the value collected in the $PL_{fb,f}$ field 256. In one embodiment of the invention, the value in the $PL_{fb,f}$ field 256 is placed into the backward priority level feedback field shown as $PL_{fb,b}$ 258. This allows the $PL_{fb,f}$ field 256 to gather priority level status information as the status cell 250 travels from the destination end-system 208 to the source end-system 206 in a manner similar to the source-to-destination status collection. Consequently, the destination end-system sets the $PL_{fb,f}$ field 256 to the lowest priority, corresponding to a priority level value of "7", and the status cell 250 is transmitted back into the network to return to the source end-system 206. During the return journey, the $PL_{fb,f}$ field 256 will again gather network load status information, this time for the connection from the destination end-system 208 to the source end-system 206. The previously collected load information stored in the $PL_{fb,b}$ 258 field will remain static for analysis at the source end-system 206.

The source end-system 206 is configured such that the failure to receive the status cell within a predetermined time period will result in an automatic decrease of the cell transfer rate. This is based on the assumption that loss of a status cell indicates that it was discarded in view of the cell transfer rate being too high, and therefore should be lowered.

In another embodiment of the invention, several pairs of $PL_{fb,f}/PL_{fb,b}$ fields can be provided in the payload 254. This can provide the user with information regarding the network load conditions with respect to various parameters, such as various time periods. For example, first, second and third pairs of $PL_{fb,f}/PL_{fb,b}$ fields can provide the network load conditions during the last 100 milliseconds, 10 seconds, and 10 minutes respectively.

Referring now to FIG. 6 and Table 1 below, an example of the calculation of the forward priority level feedback $PL_{fb,f}$ 256 and the backwards priority level feedback $PL_{fb,b}$ 258 is described.

TABLE 1

	t < 1	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6
$PL_a(\text{NODE } 202)$	5	5	5	4	4	4	4
$PL_a(\text{NODE } 204)$	—	—	—	6	6	6	6
$PL_{fb,f}$	7	7	5	5	5	5	4
$PL_{fb,b}$	—	—	—	—	5	5	5

As illustrated in FIG. 6, the status cell 222a-f is shown at times t=1 through t=6 respectively. Table 1 illustrates the $PL_{fb,f}$ 256 and $PL_{fb,b}$ 258 as compared to the allowable priority level PL_a at nodes 202 and 204 at times t<1 through t=6. At time t<1, the $PL_{fb,f}$ is initialized to the lowest priority, thereby having a preset priority level value of "7". At time t=1, the status cell 222a is transmitted from the source end-system 206 to the ATM node 202, at which time the $PL_{fb,f}$ still has the value of "7". Because the node 202 has a PL_a value of "5", $PL_{fb,f}$ 256 within status cell 222b is decremented to a value of "5" at time t=2 to reflect the current worst-case PL_a value of the connection. Node 204 has a PL_a value of "6" at time t=3, which is greater than the current state of $PL_{fb,f}$ 256 which is equal to "5". Therefore, $PL_{fb,f}$ 256 remains unchanged at time t=3, when status cell 222c exits ATM node 204.

Between times t=3 and t=4, $PL_{fb,f}$ 256 is placed into the backward priority level feedback field $PL_{fb,b}$ 258. At time t=4, status cell 222d therefore includes a $PL_{fb,b}$ field 258 that stores the value "5" which corresponds to the worst-case allowable priority level of the source-to-destination connection. Because the node 204 still has a PL_a value of "6", $PL_{fb,f}$

256 of status cell 222e remains unchanged at a value of "5" at time t=5. However, at some time between t=2 and t=3, the PL_a value at node 202 changed from "5" to "4", causing $PL_{fb,f}$ 256 in cell 222f to also be decreased to a value of "4". As can be seen from Table 1, $PL_{fb,b}$ 258 remains static during the return journey (i.e., t=4 through t=6), so that the source-to-destination $PL_{fb,f}$ can be reported to the source end-system 206.

FIG. 8 is a block diagram of an ATM node 300, representative of any of the nodes 202, 204 or additional nodes in the ATM network 200, configured for the NBR methodology. Each node may have multiple inputs from other nodes or end-stations, seen generally as links 302. The switch 304 receives each of the links 302 containing multiplexed information streams, and rearranges the information streams between input and output, as is generally known in the art. In the example of FIG. 8, switch 304 receives cell 306a at link 308, and provides the cell 306b at its output.

The ATM node 300 determines 310 whether the cell 306b is a standard data cell or a status cell configured for the NBR methodology. In one embodiment, this is accomplished by comparing a known value to a payload type value in the payload type (PT) field 260 in the header 252. Where the cell 306b is not an NBR status cell, it is a standard ATM information cell which is transferred to the cell scheduling and buffering circuitry 312, generally described in connection with FIGS. 4 and 5, which accepts and discards the cells according to the priority level of the cell and the current buffer occupancy level. Where the cell 306b is an NBR status cell, the $PL_{fb,f}$ field 256 is appropriately set 314 in the cell 306c in accordance with the current allowed priority level PL_a . The various PL_a values for each of the links coupled to the ATM node 300 are stored in a storage table 316. The $PL_{fb,f}$ field 256 is set 314 to a value equivalent to the PL_a value in table 316 when PL_a is less than the value currently residing in the $PL_{fb,f}$ field 256. Otherwise, the $PL_{fb,f}$ field 256 remains unchanged. Whether the $PL_{fb,f}$ field 256 is modified or not, the status cell 306d is provided to the cell scheduling and buffering circuitry 312 to be filtered and buffered like any standard ATM cell. Cells exit the node 300 at output links 318, where the status cell 306e of this example is shown exiting node 300 via link 320. In one embodiment, a cell scheduling and buffering block 312 is provided for each output of the node, such that each cell scheduling and buffering block 312 operates independently of other cell scheduling and buffering blocks (not shown) at the node.

In FIG. 9, there is illustrated in flow diagram form, in accordance with one embodiment of the invention, a general methodology by which NBR status cells can be used to provide feedback to a source end-station so that cell transfer rates may be optimized. A status cell is transmitted 400 from the source end-system 206 to the destination end-system 208. Intermediary nodes, such as nodes 202 and 204 of FIG. 6, detect 402 status cells sent by users at the source. The allowable priority level PL_a is stored 404 as $PL_{fb,n}$ at each of the intermediary nodes. The highest allowable priority level identified at all of the nodes traversed prior to the current node is available in the $PL_{fb,f}$ field 256 of the status cell, which can then be compared 406 to $PL_{fb,n}$ at each node. Where $PL_{fb,n} > PL_{fb,f}$ the node does not change 408 the value in the $PL_{fb,f}$ field 256. Where $PL_{fb,n} < PL_{fb,f}$ the value in the $PL_{fb,f}$ field 256 is set 410 to reflect the lowest allowable priority level identified up to this point in the connection, and is therefore decremented to equal $PL_{fb,n}$ at the node.

The status cell may encounter more nodes 412 as it progresses towards its destination. Where more intermediary

nodes are in the path of the status cell, each of the intermediary nodes must detect 402 the status cell, and the $PL_{pb,f}$ field 256 is set accordingly 404, 406, 408, 410. Where no more intermediary nodes are encountered from source to destination, the status cell, and the current load indication PL_{pb} is returned 414 to the source end system. In one embodiment of the invention, the value in the $PL_{pb,f}$ field 256 is placed into the backward priority level feedback field $PL_{pb,b}$ 258 prior to the cell's departure from the destination end-system 208. This allows new network load information to be gathered in the $PL_{pb,f}$ field 256 on the return connection without corrupting the network load information gathered on the source-to-destination connection.

When the source end-system 206 receives the returned status information, it can modify its current traffic parameters to optimize cell transmission. Where the returned status indicates that the cell traffic is relatively low 416, the user may decrease 418 the priority of newly issued ATM data cells to reflect a priority level that will likely be accepted at each of the nodes of the connection. Similarly, where the cell traffic appears to be relatively high 420, the user may increase 422 the priority of newly issued data cells. This allows adjustment of the confidence level that data cells will not be discarded at any node along the connection. The returned status is therefore used to allow the user to adjust the cell transfer rate (CTR) of cells exiting the source end-system 206 to the network.

FIGS. 10-13 illustrate various embodiments for determining the priority level feedback (PL_{pb}) in a switch in accordance with the present invention. Referring first to FIG. 10, a first embodiment for determining PL_{pb} is illustrated. In this embodiment, the PL_a values calculated at the ATM node as cells pass through the switch are monitored over a predetermined observation time period t_0 . The minimum PL_a value occurring during this predetermined time period t_0 is recognized as the lowest accepted priority level PL_a (i.e., the highest priority) for this period, thereby indicating the threshold priority level required to ensure acceptance of cells for that period.

In one embodiment, the period over which the minimum PL_a value is monitored is calculated by monitoring the number of cells entering the ATM node, and selecting the minimum PL_a from a "window" of ATM cells corresponding to a predetermined number of cells entering the ATM node. The window of cells has a size l_0 determined by:

$$l_0 = c(t_0) \quad [2]$$

where

c is the rate of cells arriving at the ATM node, and t_0 is the observation time period.

The smallest accepted priority level PL_a value inside this window of size l_0 , represented by dimension line 500, is compared with the $PL_{pb,f}$ value carried in each cell. In the example of FIG. 10, a PL_a of "5" at time slot 502 is the minimum PL_a during the window l_0 . The PL_{pb} value is therefore determined by the following mathematical equation:

$$PL_{pb} = \text{Min}(PL_a) \quad [3]$$

where $i=1, 2, \dots, l_0$.

The size l_0 of the window may be set to any desired value to provide varying degrees of accuracy. Therefore, while l_0 may be set as low as one, a preferred embodiment of the invention sets the window size l_0 to approximately correspond to the time delay of a round-trip cell path between the source and the destination. However, the observed time

interval t_0 may be set to a much larger time if providing long-term statistical priority feedback information to the source is desired.

Referring now to FIG. 11, another embodiment for determining PL_{pb} is illustrated. In this embodiment, the PL_a values calculated at the ATM node as cells pass through the switch are monitored over a predetermined observation time period t_0 , and the minimum PL_a value which occurs a predetermined number of times during this time period t_0 is recognized as the lowest accepted priority level PL_a . This method may be beneficial where single, unique events cause the minimum PL_a to inaccurately represent the typical minimum priority level at that ATM node during the period in question. For example, where the PL_a during a window l_0 500 includes values 5 and 6, but one of the PL_a values is 2, the connection may undeservingly be classified as a connection typically requiring a very high priority to avoid being discarded at an intermediary node. Therefore, FIG. 11 represents a manner of increasing the probability of providing a more typical minimum PL_a .

As illustrated in FIG. 11, a predetermined number "M" of PL_a values must be recognized before modifying the $PL_{pb,f}$ value. M equals three in FIG. 11, and PL_a is equal to 6 three times in the window l_0 500, as represented by time slots 510, 512 and 514. The PL_a value of 5 at time slot 516 was not selected because it only occurred once during the predetermined time period t_0 . The PL_a value of 7 at time slots 518, 520, 522 and 524 were not selected because it was not the minimum accepted priority level value occurring at least M times.

FIG. 12 illustrates still another embodiment for determining PL_{pb} . PL_{pb} may be calculated by averaging the PL_a values over a time period t_0 . This method is most beneficial where there are relatively few large variations in PL_a over relatively long time periods. The PL_{pb} value can therefore be determined by the following mathematical equation:

$$PL_{pb} = \text{Int} \left(\sum_{i=1}^{l_0} \frac{1}{l_0} PL_a^i \right) \quad [4]$$

where

l_0 is the number of cells inside the window, and

Int represents the nearest integer value of the resulting average.

Referring again to FIG. 12, the average of the PL_a values is determined by applying Equation [4], where the PL_a values in the 7 (l_0) times slots 530, 532, 534, 536, 538, 540 and 542 of the window are averaged over the time t_0 corresponding to the window of l_0 cells. In the example of FIG. 12, the calculated average is approximately 6.3, resulting in a $PL_{pb,f}$ of 6.

FIG. 13 illustrates yet another embodiment for determining PL_{pb} . PL_{pb} may be determined by calculating a weighted average of PL_a over a time period t_0 . The PL_a values at the remote end of the window (corresponding to the least recently occurring ATM cells in the window) are afforded less weight in the calculation as the more recent PL_a values. More emphasis is placed on the more recent PL_a values, as it may provide a better indication of the trend under changing conditions. Mathematically, this is represented by:

$$PL_{pb} = \text{Int} \left(\sum_{i=0}^{l_0} w_i PL_a^i \right) \quad [5]$$

-continued

where the total sum of weights is $\sum_i w_i = 1$.

As seen in FIG. 13, the weighted average of the PL_a values is determined by applying Equation [5], where the PL_a values in the 7 times slots 550, 552, 554, 556, 558, 560 and 562 of the window are divided into four "weighted" segments. The first weighted segment 564 includes the PL_a values corresponding to the least recently occurring cells 550 and 552 which, in this example, are provided a weight value of 0.05. The second weighted segment 566 includes the PL_a values corresponding to cells 554 and 556 which, in this example, are provided a weight value of 0.1. The third weighted segment 568 includes the PL_a values corresponding to cells 558 and 560 which are provided a weight value of 0.2 in this example. Finally, a fourth segment 570 includes the PL_a values corresponding to cell 562 which has a weight value of 0.3. Applying Equation [5] to the example of FIG. 13, the calculated weighted average results in a PL_{avg} of 6. As will be appreciated by those skilled in the art from the above description, different quantities of weighted segments may be selected, as may different quantities of PL_a values per weighted segment, without departing from the scope and spirit of the invention.

For purposes of illustration, and not of limitation, examples are provided below to illustrate the relationship between the quality of service of an NBR or SIMA connection and throughput at different priority levels. The following examples demonstrate, in accordance with one embodiment of the present invention, the relative differences in QoS with respect to adjacent priorities, such as the QoS associated with $PL=4$ in comparison to that of $PL=3$. It is noted that a higher cost or fee is typically assessed to users who request a higher priority for cells transmitted from the user's user/network interface. By way of example, the user fee may be doubled if the user wants to obtain one degree of higher priority for every cell without changing the actual bit rate. The resulting QoS of the connection, therefore, should be improved such that at least some users are willing to pay the additional charge.

EXAMPLE #1

In accordance with this illustrative example, the following assumptions and consideration are given. It is assumed that there are many identical traffic sources which generate traffic independent of the current or previous load conditions in the network. The following traffic parameters are assumed: the link capacity is $C=1$, which is useful in the examples as a means of normalization; the peak bit rate $MBR_{MAX}=0.1$, which represents 10 percent of the link capacity, C ; the ON probability at the burst (or packet) scale=0.2; and the average burst duration=1,000 time slots (i.e., the average packet size=100 cells). In addition, it is assumed that there is an upper ON/OFF layer, and that both the average ON-period and OFF-period of this layer are 100,000 time slots. The real time buffer 93 contains 200 cell locations and the non-real-time buffer 94 contains 5,000 cell locations. It is noted that the upper ON/OFF layer attempts to model the traffic process of connections, in which the determination of the number of connections is understood in the art to constitute a stochastic, random process. For example, if it is assumed that the total number of customers is represented by the variable x , then the average number of connections is $x/2$. More particularly, the number of connections is understood to be binomially distributed. As such, 100,000 time slots

represent the average holding time of a connection, and, also, the average idle period realizable by the user. As a consequence, a user is transmitting cells only if a connection is active both at the connection layer and the packet layer. A time scale parameter, α , can be obtained for the real-time and non-real-time connections:

$$\alpha_{rt}=0.025$$

$$\alpha_{nrt}=0.001$$

In this example, eight different connection types are assumed: four connections are real-time connections and four are non-real-time connections. Also, four different NBR values, which have been normalized with respect the link capacity of $C=1$, are assumed as: 0.2, 0.1, 0.05 and 0.025. The priorities corresponding to these NBR values, with $MBR_{MAX}=0.1$, are 3, 4, 5 and 6, respectively. It should be noted, however, that not all cells are assigned these exact priorities, and that especially with non-real-time connections, many cells obtain better priority values because of the affects of the averaging measuring principle. The distribution of cells having different priority levels, represented as percentages, is presented below in Table 2:

TABLE 2

PRIORITY LEVEL	REAL (SIMULATED) PERCENTAGE OF OFFERED CELLS	PERCENTAGE BASED ON PEAK RATES
1	6.1	0
2	7.9	0
3	24.3	25
4	23.5	25
5	21.5	25
6	16.8	25

In FIG. 14, there is shown a graph illustrating the relationship of average cell loss ratio, P_{loss} , as a function of priority level for four specific load levels, r . In particular, line-800 represents an overall average load level of 0.72 for 9 connections of each connection type (i.e., real-time and non-real-time connection types). Line-802 depicts an average load level of 0.80 for 10 connections of each connection type. Further, line-804 represents an average load level of 0.88 for 11 connections of each connection type, and line-806 represents an average load level of 0.96 for 12 connections of each connection type. It is noted that, in the case of line-802 indicating a load level of 0.80, the cell loss ratios, P_{loss} , for real-time and non-real-time cells are indicated by dotted and broken lines, respectively.

Given, for example, a traffic scenario where the operator wants to offer a cell loss ratio of 10^{-6} to cells with priority 4, the total load can be approximately 0.75. It can be assumed that this average cell loss ratio is sufficient for most video applications. Give the same traffic load conditions, priority level 5, which corresponds to $P_{loss}=10^{-4}$, can meet the requirements of many voice applications, while priority 6, which corresponds to $P_{loss}=3 \cdot 10^{-3}$, is suitable for a TCP/IP type of file transfer, provided that there is an adequate packet discarding scheme in place.

It should be emphasized, however, that the difference in cell loss ratio between adjacent priorities depends strongly on the offered traffic process and, in particular, the inherent control loops of the NBR or SIMA service. When the user perceives an unsatisfactory QoS, for example, the user can, and should, change either the actual bit rate or the nominal bit rate of the connection. In either case, the priority distribution changes as well. Nevertheless, if this phenomenon is temporarily ignored, the basic behavior of priority distribution may be further appreciated by making the following

simplifying assumption: If it is assumed that all traffic variations are slow as compared to the measuring period and buffer size, then a well-known, conventional ATM approach to approximating cell loss ratio may be used, with the additional requirement that the eight NBR priority levels are taken into account.

If the loss ratio of cells with priority k is denoted by $P_{loss,k}$ and the average loss ratio of cells with a priority of 0 to k is denoted by $P_{loss,k}^*$, then the following equation, which ignores buffering effect, provides that:

$$P_{loss,k} = \frac{\sum_{j:\lambda_j^* > c} Pr(\lambda_i^* = \lambda_j)(\lambda_j - c)}{\rho_k^* c} \quad [6]$$

$$P_{loss,0} = P_{loss,0}^*$$

$$P_{loss,k} = \frac{\rho_k^* P_{loss,k} - \rho_{k-1}^* P_{loss,k-1}}{\rho_k^* - \rho_{k-1}^*} \text{ for } k = 1 \dots 7$$

where, λ_k^* represents the momentary bit rate level of all cells with a priority of 0 to k , ρ_k^* represents the average offered load produced by these cells, and c represents the link capacity. The probability $Pr\{\lambda_k^* = \lambda_j\}$ can be calculated in a straightforward manner using known convolution techniques.

EXAMPLE #2

For purposes of further illustration, a second example is provided which assumes the same sources described in Example #1, except for the long ON and OFF periods. Because of the long periods reflected in Example #1, the peak rate always determines the cell priority. As the buffers are typically not capable of filtering any traffic variations, the allowed load in Example #2 is much lower than that in the original case of Example #1.

In FIG. 15, there is illustrated in graphical form a relationship between cell loss ratio as a function of priority level for different load levels, r . It is assumed in FIG. 15 that the peak cell rate of each connections depicted by solid lines 820, 822, 824 is 0.1, the peak cell rate of each connection depicted by the broken line-826 is 0.2, and the peak cell rate of each connection depicted by the dotted line-828 is 0.05.

FIG. 15 shows the cell loss probabilities obtained by application of Equation [6] for different priorities, depicted by the three solid lines, line-820, line-822, and line-824. In addition, two slightly different traffic cases are represented by the dotted line-828 and the broken line-826. The effects of changing the traffic variations are reflected in the graph provided in FIG. 15. The actual change in traffic variations is a direct consequence of doubling or halving bit rates and NBR values.

In a network that embraces the NBR/SIMA service concept, an increase of traffic variations has two main effects if the operator keeps the QoS of priority level 4 unchanged. First, the allowed load level is decreased in the same way as in conventional ATM, and second, the difference in cell loss ratio between adjacent priority level decreases. For purposes of providing a rough estimate of QoS based on FIGS. 14 and 15, it may be assumed that if priority level 4 offers a cell loss probability of 10^{-6} , then the cell loss probability will be approximately 10^{-4} to 10^{-3} with priority level 5 depending on the overall traffic variations. The cell loss ratio with priority level 3 may be assumed to be less than 10^{-9} , unless the traffic variations are very pronounced.

Although the above examples provide illustrations of relationships between QoS and priority levels, it may be unfruitful to attempt to exactly determine the allowed load

or the cell loss difference between adjacent priority levels until user reactions to different QoS and usage charges are evaluated. In an NBR/SIMA service environment, a schedule of charges based on different QoS levels may be determined, in a certain sense, automatically. For example, if the difference in cell loss ratio between priority levels 4 and 5 is very small, it can be assumed that some of the connections will tend to move from priority level 4 to level 5 because of a lower assessed charge. This change indicates, apparently, that the cell loss ratio of priority level 4 decreases and the cell loss ratio of priority level 5 increases. It can be reasonably assumed that this type of movement continues until the QoS difference corresponds to the average user's expectation of a reasonable charging structure.

Similar concerns are raised with regard to the differences in charging which occur automatically during busy hours in contrast to idle hours. For example, it would appear reasonable to charge higher prices during periods of high load for a certain QoS, and charge lower prices during low load periods. However, it is believed desirable to avoid instigating a charging policy that assesses different fees for a certain NBR during busy and idle periods, which also avoids increasing the complexity of the charging scheme. The naturally occurring "supply and demand" effect may tend to automatically even out the load between busy and idle hours. It is anticipated that, if a user is displeased with the observable difference in QoS during busy and idle periods, the user may be motivated to pay a different rate during such periods.

EXAMPLE #3

Another important transmission control aspect is the modification and adjustment of the cell transfer rate, or similarly, the measured bit rate (MBR). It is possible to provide those sources, which are able to change their bit rate, information about the current load condition on their respective connections. This example provides a performance outlook for these connections.

For the present example, the background traffic process generated by the sources previously described in Examples 1 and 2 is used. It is also assumed that there are 10 connections of each type, thereby providing an average background load of $r=0.080$. There are also three sources (hereinafter referred to as feedback sources FBS1, FBS2 and FBS3) that adjust their transmission rate according to the feedback information received via the network status cells. All of these feedback sources are assumed to have an NBR of 0.01.

The feedback sources are similar to each other, except for their time periods which are used in determining PL_n information. More specifically, the feedback source time periods are 10,000 time slots for FBS1, 30,000 time slots for FBS2, and 100,000 time slots for FBS3. For comparison purposes to the feedback sources, three constant bit rate connections (source connections C4, C5 and C6) are also discussed, which have the following parameters (link capacity=1):

TABLE 3

	Bit-rate/NBR	NBR	Priority Level (PL)
C4	1.4	0.01	4
C5	2.8	0.01	5
C6	5.5	0.01	6

The cell sources in Table 3 are set such that they transmit at a rate that is slightly less than the threshold that borders the next higher priority level. For example, using Equation 1, bit rate/NBR ratio equal to 1.5 would result in a cell

priority level of 5 (integer portion of 5.08), while a bit rate/NBR ratio of 1.4 yields a cell priority level of 4 (integer portion of 4.98). The feedback sources have been set to the same bit rate values in this example in order to optimally utilize the network capacity.

In FIG. 16, there is shown a graph illustrating the results of a simulation comparing the constant bit rate connections to those using feedback sources. The cell loss ratio is plotted on the vertical axis 850 as a function of accepted bit rate/NBR plotted on the horizontal axis 852. The graph illustrates that constant bit rate sources, C4 854, C5 856, and C6 858, obtain better loss vs. throughput characteristics when they are transmitting cells at lower priority level values (i.e., at higher priorities). However, the results indicate that feedback sources, which adjust their rate using feedback information from the status cells, are useful when the background traffic is changing slowly. The feedback sources FB1 870, FB2 872 and FB3 874 can adapt to changes, while constant bit rate sources cannot take advantage of the changing load of the network. With fast changing background traffic, feedback sources are not able to adjust to changes fast enough, resulting in an increased cell loss ratio.

FIG. 16 also illustrates the use of one source that generally performs as a kind of TCP source 876 with NBR=0.01. This source halves its transmission rate when it receives information about a lost cell, and increases its transmission rate by 10% if it does not receive information regarding lost cells over a time period of 10,000 time slots. Comparison reveals that this kind of source loses more cells than corresponding feedback sources. This is expected, as this type of source has a slower reaction to changes.

Another aspect to be considered involves the ability of the connections to be adjusted in during periods of sudden capacity changes, as well as how the capacity division between different feedback connections will be managed. Referring now to FIG. 17, a graphical illustration of the throughput/capacity (vertical axis 900) from each of the sources is shown as a function of time (horizontal axis 902). FIG. 17 provides an example of four feedback sources (FB1 904, FB2 906, FB3 908 and FB4 910) are transmitting cells to a network node, where FB1 904 and FB2 906 have an NBR=0.25, and FB3 908 and FB4 910 have an NBR=0.0625. At a time corresponding to 30,000 time slots, a uniform source begins to transmit, shown as step function 912, with a cell rate of PCR=0.333 at an NBR=0.333. At a time corresponding to 60,000 time slots, the uniform source terminates its cell transmission.

As can be seen in FIG. 17, the feedback sources (FB) are able to adjust their transmission rates in view of the load step function 912 on the connection. All the feedback sources have roughly the same decrease of throughput when the step source 912 is switched on. After the step source is switched off, the feedback sources restore their original throughput.

It will, of course, be understood that various modifications and additions can be made to the preferred embodiments discussed hereinabove without departing from the scope or spirit of the present invention. Accordingly, the scope of the present invention should not be limited by the particular embodiments discussed above, but should be defined only by the claims set forth below and equivalents thereof.

What is claimed is:

1. A method of acquiring load information for a network connection defined between a source unit and a destination unit, comprising:

determining a threshold priority level for each network node defined along the network connection, the thresh-

old priority level being used when accepting or discarding information elements received by each node; computing a worst-case threshold priority level using the threshold priority levels determined for the nodes; transmitting a status cell along the network connection: maintaining a record of the worst-case threshold priority level by storing in the status cell information indicative of the worst-case threshold priority level; and providing the worst-case threshold priority level to the source unit.

2. The method of claim 1, wherein computing the worst-case threshold priority level comprises determining a minimum threshold priority level of the threshold priority levels determined for the nodes.

3. The method of claim 1, wherein determining the worst-case threshold priority level comprises determining a maximum threshold priority level of the threshold priority levels determined for the nodes.

4. The method of claim 1, wherein maintaining the record of the worst-case threshold priority level further comprises: storing in a first field of the status cell a worst-case threshold priority level computed by the nodes as the status cell progresses along the connection from the source unit to the destination unit; and

storing in a second field of the status cell a worst-case threshold priority level computed by the nodes as the status cell progresses along the connection from the destination unit to the source unit.

5. The method of claim 1, wherein determining the threshold priority level for each network node defined along the network connection further comprises determining a priority level for each information element transmitted over the connection using a ratio of a nominal bit rate established for the connection and an actual bit rate of the connection, the priority level indicating a probability of information element acceptance at each network node.

6. The method of claim 1, wherein the network connection constitutes a nominal bit rate connection.

7. A method of adjusting a transmission rate of information elements transmitted between a source unit and a destination unit over a network connection, comprising:

transmitting the information elements and a status cell over the network connection;

determining a threshold priority level for each network node defined along the network connection, the threshold priority level being used when accepting or discarding information elements received by each node;

storing, in the status cell, information indicative of a worst-case threshold priority level using the threshold priority values determined for each of the nodes; and adjusting the transmission rate of the information elements in response to the worst-case threshold priority level information stored in the status cell.

8. The method of claim 7, wherein:

transmitting the information elements further comprises assigning a priority level to each information element transmitted from the source unit, the priority level indicating a probability of information element acceptance at each node; and

adjusting the transmission rate of the information elements further comprises adjusting the priority level of each information element transmitted from the source unit in response to the worst-case threshold priority level information stored in the status element.

9. The method of claim 7, wherein storing the worst-case threshold priority level information further comprises:

storing in a first field of the status element a worst-case threshold priority level as the status element progresses along the connection from the source unit to the destination unit; and

storing in a second field of the status element a worst-case threshold priority level as the status element progresses along the connection from the destination unit to the source unit.

10. The method of claim 7, wherein determining the threshold priority level for each network node defined along the network connection further comprises computing a priority level for each information element transmitted over the connection using a ratio of a nominal bit rate established for the connection and an actual bit rate of the connection, the priority level indicating a probability of information element acceptance at each network node.

11. The method of claim 7, wherein storing the worst-case threshold priority level information further comprises:

comparing a threshold priority level of a node receiving the status element to a threshold priority level of a node transmitting the status element; and

computing the worst-case threshold priority level using the respective threshold priority levels of the transmitting and receiving nodes.

12. The method of claim 11, wherein computing the worst-case threshold priority level further comprises computing either one of a maximum or a minimum of the worst-case threshold priority levels of the respective transmitting and receiving nodes.

13. The method of claim 7, wherein the network connection constitutes a nominal bit rate connection.

14. A system for adjusting a transfer rate of information cells transmitted over a network connection, comprising:

a plurality of network nodes defined along the network connection;

a source unit that transmits the information cells and a status cell over the network connection, the source unit establishing a priority level for each of the information cells indicative of a probability of information element acceptance at a network node;

a destination unit that receives the information cells and the status cell transmitted over the network connection, and transmits the status cell back to the source unit; and

a processor, provided at each node, that computes a threshold priority level used in accepting or discarding information cells received by the node, and computes a worst-case threshold priority level using the computed threshold priority level and threshold priority level information stored in the status cell;

wherein the source unit effects adjustment of priority levels for information cells subsequently output from the source unit in response to the worst-case threshold priority level information stored in the status cell.

15. The system of claim 14, wherein:

the processor of each node stores a worst-case threshold priority level computed by the node in a first field of the status cell as the status cell progresses along the connection from the source unit to the destination unit; and the processor of each node stores a worst-case threshold priority level computed by the node in a second field of the status cell as the status cell progresses along the connection from the destination unit back to the source unit.

16. The system of claim 14, wherein the source unit establishes the priority level for each information element transmitted over the connection using a ratio of a nominal bit rate established for the connection and an actual bit rate of the connection.

17. The system of claim 14, wherein:

the threshold priority level information stored in the status cell represents a worst-case threshold priority level computed by nodes receiving the status cell prior to a particular node receiving the status cell; and

the processor of the particular node computes the worst-case threshold priority level by determining either one of a maximum or a minimum of the threshold priority level computed by the particular node and the threshold priority level information stored in the status cell.

18. The system of claim 14, wherein the network connection constitutes a nominal bit rate connection.

19. A method of acquiring load information for a network connection defined between a source unit and a destination unit, comprising:

determining a threshold priority level for each network node defined along the network connection, the threshold priority level being used when accepting or discarding information elements received by each node;

computing a worst-case threshold priority level using the threshold priority levels determined for the nodes;

maintaining a record of a forward worst-case threshold priority level for the connection defined from the source unit to the destination unit;

maintaining a record of a backward worst-case threshold priority level for the connection defined from the destination unit to the source unit; and

providing the forward and backward worst-case threshold priority levels to the source unit.

* * * * *